

Energy and Performance Management of Green Data Centers: A Profit Maximization Approach

Mahdi Ghamkhari, *Student Member, IEEE*, and Hamed Mohsenian-Rad, *Member, IEEE*

Abstract—While a large body of work has recently focused on reducing data center’s energy expenses, there exists no prior work on investigating the *trade-off* between minimizing data center’s energy expenditure and maximizing their revenue for various Internet and cloud computing services that they may offer. In this paper, we seek to tackle this shortcoming by proposing a systematic approach to maximize green data center’s profit, i.e., revenue minus cost. In this regard, we explicitly take into account practical *service-level agreements* (SLAs) that currently exist between data centers and their customers. Our model also incorporates various other factors such as availability of local renewable power generation at data centers and the stochastic nature of data centers’ workload. Furthermore, we propose a novel *optimization-based* profit maximization strategy for data centers for two different cases, without and with behind-the-meter renewable generators. We show that the formulated optimization problems in both cases are *convex programs*; therefore, they are tractable and appropriate for practical implementation. Using various experimental data and via computer simulations, we assess the performance of the proposed optimization-based profit maximization strategy and show that it significantly outperforms two comparable energy and performance management algorithms that are recently proposed in the literature.

Keywords: Green data centers, behind-the-meter renewable power generation, energy and performance management, service-level agreements, profit maximization, convex optimization.

I. INTRODUCTION

The growing demand for Internet services and cloud computing has significantly increased the electric power usage associated with large data centers - such as those owned and operated by Google, Microsoft, and Amazon - over the past few years. Each data center includes hundreds of thousands of computer servers, cooling equipment, and substation power transformers. For example, consider Microsoft’s data center in Quincy, WA. It has 43,600 square meters of space and uses 4.8 kilometers of chiller piping, 965 kilometers of electric wire, 92,900 square meters of drywall, and 1.5 metric tons of backup batteries. The peak power consumption of this facility is 48 megawatts, which is enough to power 40,000 homes [1]. As another example, the National Security Agency is currently building a massive data center at Fort Williams in Utah which is expected to consume over 70 megawatts electricity [2].

Due to the increasing cost of electricity associated with data centers, there has been a growing interest towards developing techniques and algorithms to minimize data centers’ energy expenditure. One thread of research focuses on reducing the amount of energy consumed by computer servers [3]. Another

thread of research is *dynamic cluster server configuration* to reduce the total power consumption by consolidating workload only on a subset of servers and turning off the rest, during low workload hours [4], [5]. A similar approach is *dynamic CPU clock frequency scaling* [6], [7]. In this approach, a higher frequency, imposing higher energy consumption, is chosen only at peak workload hours. Finally, some recent studies aimed to utilize *price-diversity* in deregulated electricity markets as well as *locational-diversity* in renewable power generation. The idea is to constantly monitor the price of electricity and the amount of renewable power generated at different regions and forward the workload towards data centers that are located in regions with the lowest electricity price [8], [9] or highest renewable power available [10], [11].

While a large body of work has addressed minimizing data centers’ cost, e.g., in [4]–[11], to the best of our knowledge, no prior work has addressed the *trade-off* between *minimizing data center’s energy expenditure* and *maximizing their revenue* for various Internet and cloud computing services that they offer. Such trade-off is due to the fact that minimizing data center’s energy cost is achieved essentially by turning off some servers, scaling down CPU clocks, or migrating some workload, which can all potentially lead to degrading the quality-of-services offered by data center and consequently its income, considering the stochastic nature of workload. Therefore, in this paper, we seek to tackle this shortcoming by proposing a systematic approach to *maximize green data center’s profit*, i.e., revenue minus cost. In this regard, we explicitly take into account practical *service-level agreements* (SLAs) that currently exist between data centers and their customers. In summary, our contributions are as follows:

- We develop a mathematical model to capture the trade-off between minimizing a data center’s energy cost versus maximizing the revenue it receives for offering Internet services. We take into account computer server’s power consumption profiles, data center’s power usage effectiveness, price of electricity, availability of renewable generation, total workload in terms of the rate at which service requests are received at each time of day, practical service-level agreements and their parameters for service deadline, service payment, and service violation penalty.
- We propose a novel *optimization-based profit maximization* strategy for data centers for two different cases, without and with behind-the-meter renewable generators. The latter is the scenario applicable to green data centers. We show that the formulated optimization problems in both cases are convex programs; therefore, they are tractable and appropriate for practical implementation.
- We use *experimental data*, e.g., for workload, price of

Manuscript received April 20, 2012; revised August 15, 2012 and November 25, 2012; accepted December 29, 2012. This work was supported in part by the National Science Foundation through grant ECCS 1253516.

The authors are with the EE Department, University of California at Riverside, Riverside, CA, USA, e-mails: {ghamkhari, hamed}@ee.ucr.edu.

electricity, renewable power generation, and SLA parameters, to assess the accuracy of the proposed mathematical model for profit and also the performance of the proposed optimization-based profit maximization strategy via computer simulations. We show that our proposed optimization-based designs significantly outperform two comparable energy and performance management algorithms that have recently been proposed in the literature.

The rest of this paper is organized as follows. The system model and notations are defined in Section II. The proposed optimization-based profit maximization strategy is presented in Section III. Simulation results are presented in Section IV. Conclusions and future work are discussed in Section V.

II. SYSTEM MODEL

Consider an Internet or cloud computing data center with M_{max} computer servers as shown in Fig. 1. Next, we explain the system model in terms of power consumption, price of electricity, incoming workload, and quality-of-service.

A. Power Consumption

The total amount of power consumption in a data center is obtained by adding the total power consumption at the computer servers to the total power consumption at the facility, e.g., for cooling, lighting, etc. For a data center, *power usage effectiveness* (PUE), denoted by E_{usage} , is defined as the ratio of the data center's total power consumption to the data center's power consumption at the computer servers [12]. The PUE is considered as a measure for data center's energy efficiency. Currently, the typical value for most enterprise data centers is 2.0 or more. However, recent studies have suggested that many data centers can soon reach a PUE of 1.7. A few state-of-the art facilities have reached a PUE of 1.2 [12].

Let P_{idle} denote the average idle power draw of a single server and P_{peak} denote the average peak power when a server is handling a service request. The ratio P_{peak}/P_{idle} denotes the power elasticity of servers. Higher elasticity means less power consumption when the server is idle, not handling any service request. Let $M \leq M_{max}$ denote the number of servers that are 'on' at data center. The total electric power consumption associated with the data center can be obtained as [13]–[16]:

$$P = M[P_{idle} + (E_{usage} - 1)P_{peak} + (P_{peak} - P_{idle})U], \quad (1)$$

where U is the CPU utilization of servers. From (1), the power consumption at data center increases as we turn on more computer servers or run servers at higher utilization.

B. Electricity Price

The electricity pricing models that are deployed for each region usually depend of whether the electricity market is regulated or deregulated in that region. The electricity prices often have flat rates and do not change during the day when the electricity market is regulated. On the other hand, the prices may significantly vary during the day when the electricity market is deregulated as the prices would reflect

the fluctuations in the *wholesale electricity market*. Some of the non-flat pricing tariffs in deregulated electricity markets include: *Day-ahead pricing* (DAP), *time-of-use pricing* (TOUP), *critical-peak pricing* (CPP), and *real-time pricing* (RTP). Our proposed energy and performance management design in this paper is applicable to not only flat rate but also non-flat rate pricing tariffs. In our system model, the instantaneous price of electricity is denoted by ω . In Section III, we will use pricing information to obtain data center's cost of electricity.

C. Renewable Power Generation

In order to reduce cost of electricity, a data center may be equipped with behind-the-meter renewable generators, e.g., a wind turbine, in addition to being connected to the power grid. Let G denote the renewable power generated by renewable generators. The amount of power exchange with the power grid is obtained as $P - G$. If local renewable power generation is lower than local power consumption, i.e., $P > G$, then $P - G$ is positive and the power flow is in the direction from the power grid to the data center. If $P = G$ then the data center operates as a *zero-net energy facility* [17]. If $P < G$, then $P - G$ is negative and the power flow is in the direction from the data center to the power grid. However, in this case, whether the grid compensates the data center's injected power depends on the market model being used. Although, in some areas, the utility pays for the power injected into the grid, currently, there is no specific market model for behind-the-meter generation in most regions in the U.S. Therefore, in our system model, while we allow a data center to inject its excessive renewable generation into the power grid, we assume that it does not receive compensation for the injected power.

D. Quality-of-Service

Because of the limited computation capacity of data centers and given the stochastic nature of most practical workload, data centers cannot process the incoming service requests, immediately after they arrive. Therefore, all arriving service requests are first placed in a queue until they can be handled by any available computer. In order to satisfy quality-of-service requirements, the waiting time/queuing delay for each incoming service request should be limited within a certain range which is determined by the *Service Level Agreement* (SLA). The exact SLA depends on the type of service offered which may range from cloud-based computational tasks to video streaming and HTML web services. Examples of two typical SLAs for an Internet or cloud computing data center are shown in Fig. 2 [18]. In this figure, each SLA is identified by three non-negative parameters D , δ , and γ . Parameter D indicates the maximum waiting time that a service request can tolerate. Parameter δ indicates the service money that the data center receives when it handles a single service request before deadline D . Finally, parameter γ indicates the penalty that the data center has to pay to its customers every time it *cannot* handle a service request before deadline D . For the Gold SLA in Fig. 2, we have $D = 300$ ms, $\delta = 7 \times 10^{-5}$ dollars, and $\gamma = 3.5 \times 10^{-5}$ dollars. For the Silver SLA, we have $D = 200$ ms, $\delta = 5 \times 10^{-5}$ dollars, and $\gamma = 1.6 \times 10^{-5}$ dollars.

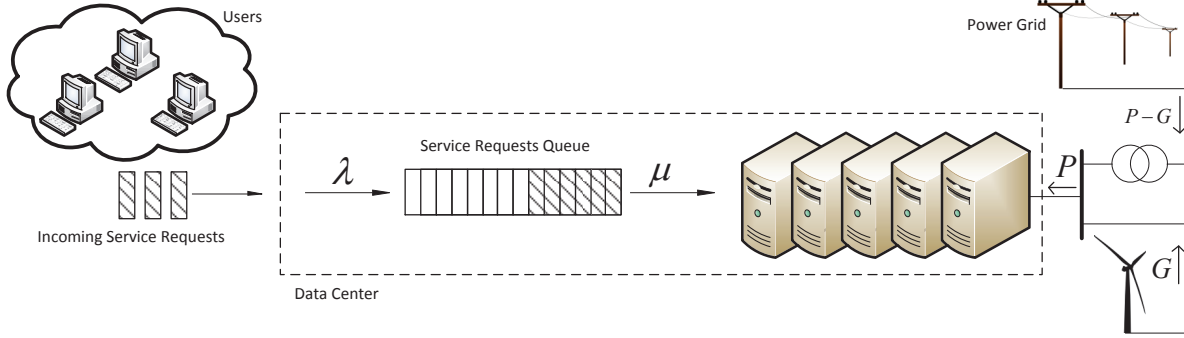


Fig. 1. In the studied system model, service requests are first placed in a queue before they are handled by one of the available computer servers.

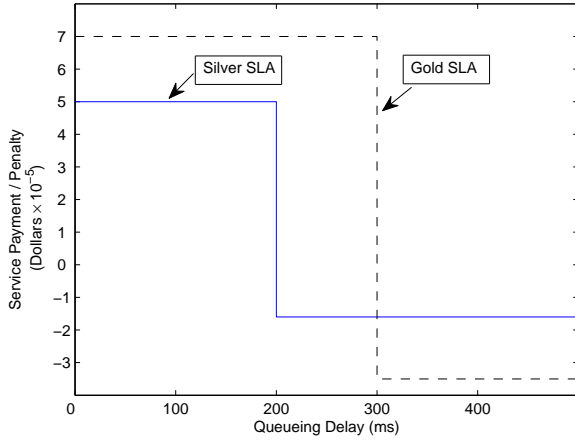


Fig. 2. Two sample service-level agreements (SLAs) in data centers.

E. Service Rate

Let μ denote the rate at which service requests are removed from the queue and handled by a server. The service rate depends on the number of servers that are switched on. Let S denote the time it takes for a server to finish handling a service request. Each server can handle $\kappa = 1/S$ service requests per second. Therefore, the total service rate is obtained as

$$\mu = \kappa M \quad \Rightarrow \quad M = \frac{\mu}{\kappa}. \quad (2)$$

As we increase the number of switched on servers and accordingly the service rate, more service requests can be handled before the SLA-deadline D , which in turn increases the payments that the data center receives as explained in Section II-D. On the other hand, it will also increase the data center's power consumption and accordingly the data center's energy expenditure as explained in Sections II-A and II-B. Therefore, there is a *trade-off* when it comes to selecting the data center's service rate, as we will discuss in detail next.

III. PROBLEM FORMULATION

The rate at which *service requests* arrive at a data center can vary over time. To improve data center's performance, the number of switched on servers M should be adjusted according to the rate of receiving service requests. More servers should be turned on when service requests are received

at higher rates. By monitoring service request rate and adjusting the number of switched on servers accordingly, only a proportional-to-demand portion of servers will be on. This results in reducing data center's power consumption.

Because of the tear-and-wear cost of switching servers on and off, and also due to the delay in changing the status of a computer, M cannot be change rapidly. It is rather desired to be updated only every few minutes. Therefore, we divide running time of data center into a sequence of time slots $\Lambda_1, \Lambda_2, \dots, \Lambda_N$, each one with length T , e.g. $T = 15$ minutes. The number of switched on servers are updated only at the beginning of each time slot. For the rest of this section, we focus on mathematically modeling the energy cost and profit of the data center of interest at each time slot $\Lambda \in \Lambda_1, \Lambda_2, \dots, \Lambda_N$ as a function of service rate μ and consequently as a function of M , based on (2).

A. Revenue Modeling

Let $q(\mu)$ denote the probability that the waiting time for a service request exceeds the SLA-deadline D . Obtaining an analytical model for $q(\mu)$ requires a queueing theoretic analysis that we will provide in Section III-C. Next, assume that λ denotes the average rate of receiving service requests within time slot Λ of length T . The total revenue collected by the data center at the time slot of interest can be calculated as

$$\text{Revenue} = (1 - q(\mu))\delta\lambda T - q(\mu)\gamma\lambda T, \quad (3)$$

where $(1 - q(\mu))\delta\lambda T$ denotes the total payment received by the data center within interval T , for the service requests that are handled before the SLA-deadline, while $q(\mu)\gamma\lambda T$ denotes the total penalty paid by the data center within interval T for the service requests that are not handled before the SLA-deadline.

B. Cost Modeling

Within time interval T , each turned on server handles

$$\frac{T(1 - q(\mu))\lambda}{M} \quad (4)$$

service requests. This makes each server busy for $T(1 - q(\mu))\lambda/\kappa M$ seconds. By dividing the total CPU busy time by T , the CPU utilization for each server is obtained as

$$U = \frac{(1 - q(\mu))\lambda}{\kappa M}. \quad (5)$$

Replacing (2) and (5) in (1), the power consumption associated with the data center at the time slot of interest is obtained as

$$P = \frac{a\mu + b\lambda(1 - q(\mu))}{\kappa}, \quad (6)$$

where

$$a \triangleq P_{idle} + (E_{usage} - 1)P_{peak}, \quad (7)$$

$$b \triangleq P_{peak} - P_{idle}. \quad (8)$$

Multiplying (6) by the electricity price ω , the total energy cost at the time interval of interest is obtained as¹

$$Cost = T\omega \left[\frac{a\mu + b\lambda(1 - q(\mu))}{\kappa} \right]. \quad (9)$$

C. Probability Model of $q(\mu)$

Consider a new service request that arrives within time slot Λ . Let Q denote the number of service requests waiting in the service queue right before the arrival of the new service request. All the Q requests must be removed from the queue, before the new request can be handled by any available server. Since the data center's service rate is μ , it takes Q/μ seconds until all existing requests are removed from the queue. Hence, the new service request can be handled after Q/μ seconds since its arrival. According to the SLA, if $Q/\mu \leq D$, then the service request is handled before the deadline D . If $Q/\mu > D$, the service request is not handled before the deadline D and it is dropped. Therefore, we can model the SLA-deadline by a *finite-size queue* with the length μD . A service request can be handled before the SLA-deadline, if and only if it enters the aforementioned finite size queue. We assume that the service request rate has an arbitrary and *general* probability distribution function. On the other hand, since the service rate μ is fixed over each time interval of length T , $q(\mu)$ can be modeled as the *loss probability* of a G/D/1 queue. Therefore, following the queuing theoretic analysis in [19], we can obtain

$$q(\mu) = \alpha(\mu) e^{-\frac{1}{2} \min_{n \geq 1} m_n(\mu)}, \quad (10)$$

where

$$\alpha(\mu) = \frac{1}{\lambda\sqrt{2\pi\sigma}} e^{\frac{(\mu-\lambda)^2}{2\sigma^2}} \int_{\mu}^{\infty} (r - \mu) e^{-\frac{(r-\lambda)^2}{2\sigma^2}} dr \quad (11)$$

and for each $n \geq 1$ we have

$$m_n(\mu) = \frac{(D\mu + n(\mu - \lambda))^2}{nC_{\lambda}(0) + 2 \sum_{l=1}^{n-1} C_{\lambda}(l)(n-l)}. \quad (12)$$

Here, $\sigma = C_{\lambda}(0)$ and C_{λ} denotes the auto-covariance of the service request rate's probability distribution function. It is known that the model in (10) is most accurate when the service request rate can be modelled as a Gaussian Process, but it also works well for any general probability distribution [19], as we will confirm through simulations in Section IV. Our model is

¹Here, since our focus is on energy and performance management of data centers, the cost model only indicates the cost of electricity. However, other cost items can also be included in the model in a similar fashion.

particularly more accurate than the existing models that are based on Poisson workload arrival distributions, e.g., in [20].

Before we end this section, we shall clarify our assumption on discarding service requests that are not handled before their deadlines. The benefits of this approach are addressed in detailed in [21]. Based on the experimental customer studies in [22], [23], the authors in [21] argue that for most data center application traffic, a network flow is useful, and contributes to application throughput and operator revenue, *if and only if* it is completed within its deadline. For example, services such as Web search, retail, advertisement, social networking and recommendation systems, while are very different, share a common underlying theme that they need to serve users in a timely fashion. Consequently, when the time expires for a service request, responses, irrespective of their completeness, are shipped out. However, the completeness of the responses directly governs their quality, and in turn, operator's revenue based on the contracted SLA, as we explained in Section III-A. Some other studies that similarly assume dropping service requests if they cannot meet the deadline include [24]–[27].

D. Profit Maximization without Local Renewable Generation

At each time slot Λ , data center's profit is obtained as

$$Profit = Revenue - Cost, \quad (13)$$

where revenue is as in (3) and cost is as in (9). We seek to choose the data center's service rate μ to maximize profit. This can be expressed as the following optimization problem:

$$\begin{aligned} \text{Maximize} \quad & T\lambda [(1 - q(\mu))\delta - q(\mu)\gamma] - \\ & T\omega \left(\frac{a\mu + b\lambda(1 - q(\mu))}{\kappa} \right), \end{aligned} \quad (14)$$

where $q(\mu)$ is as in (10) and M_{\max} denotes the total number of servers available in the data center. We note that the service rate μ is lower bounded by λ . This is necessary to assure stabilizing the service request queue [11], [19], [20]. We also note that Problem (14) needs to be solved at the beginning of every time slot $\Lambda \in \{\Lambda_1, \dots, \Lambda_N\}$, i.e., once every T minutes.

E. Profit Maximization with Local Renewable Generation

When a data center is supplied by both the power grid and also a local behind-the-meter renewable generator, then the optimum choice of service rate for maximizing profit is obtained by solving the following optimization problem

$$\begin{aligned} \text{Maximize} \quad & T\lambda [(1 - q(\mu))\delta - q(\mu)\gamma] - \\ & T\omega \left[\frac{a\mu + b\lambda(1 - q(\mu))}{\kappa} - G \right]^+, \end{aligned} \quad (15)$$

where $[x]^+ = \max(x, 0)$. Note that, $(a\mu + b\lambda(1 - q(\mu)))/\kappa - G$ indicates the amount of power to be purchased from the grid. As discussed in Section II-C, if this term is negative, the data center's electricity cost will be zero, given the assumption that the grid does not provide compensation for the injected power.

F. Solution and Convexity Properties

In this Section, we characterize optimization problems (14) and (15), and show that they can be solved using efficient optimization techniques, such as the *interior point method* (IPM) [28]. First, consider the following useful theorem.

Theorem 1: The probability function $q(\mu)$ in (10) is *non-increasing* and *convex* if the service rate is limited to

$$\mu \in \left[\lambda + \frac{\sqrt{2}}{2}\sigma, \lambda + \sqrt{\frac{16/\pi}{3 - \frac{8}{\pi} + \sqrt{9 - \frac{16}{\pi}}}}\sigma \right]. \quad (16)$$

The proof of Theorem 1 is presented in Appendix A. Note that, the interval in (16) can be approximately expressed as

$$[\lambda + 0.7071\sigma, \lambda + 1.4477\sigma]. \quad (17)$$

Next, we note that handling a single service request increases power consumption of a single server from P_{idle} to P_{peak} for $1/\kappa$ seconds. This increases the energy cost by $\omega(P_{peak} - P_{idle})/\kappa = \omega b/\kappa$ and also increases the revenue by δ . Thus, running the data center is *profitable* only if

$$\delta - \omega b/\kappa > 0. \quad (18)$$

From this, together with Theorem 1, we can now provide the following key results on tractability of Problems (14) and (15).

Theorem 2: Assume that condition (18) holds and the service rate μ is limited in the range indicated in (16). (a) Problem (14) is convex in its current form. (b) Problem (15) is equivalent to the following convex optimization problem:

$$\begin{aligned} & \text{Maximize}_{\lambda \leq \mu \leq \kappa M_{max}} && T\lambda[(1 - q(\mu))\delta + q(\mu)\gamma] - \\ & && T\omega\left(\frac{a\mu + b\lambda(1 - q(\mu))}{\kappa} - G\right) \quad (19) \\ & \text{Subject to} && a\mu + b\lambda(1 - q(\mu)) \geq G. \end{aligned}$$

The proof of Theorem 2 is given in Appendix B. Note that, two optimization problems are called *equivalent* if they have equal optimal solutions such that solving one can readily solve the other one [28, pp. 130-135]. From Theorem 2, Problems (14) and (15) are both tractable and can be solved efficiently using convex programming techniques, such as the IPM [28].

An interesting extension for the design problem in (19) is the case when an SLA requires that the probability of dropping a packet is upper bounded by a constant L . This requirement can be incorporated in our design by adding the following constraint to optimization problems (14), (15) and (19):

$$q(\mu) \leq L. \quad (20)$$

However, we can show that the above constraint is convex. Note that, from Theorem 1, $q(\mu)$ is a convex function. Therefore, constraint 20 forms a convex set [28, Section 4.2.1] and the convexity of problems (14), (19) still holds. As a result, the proposed convex optimization framework is still valid.

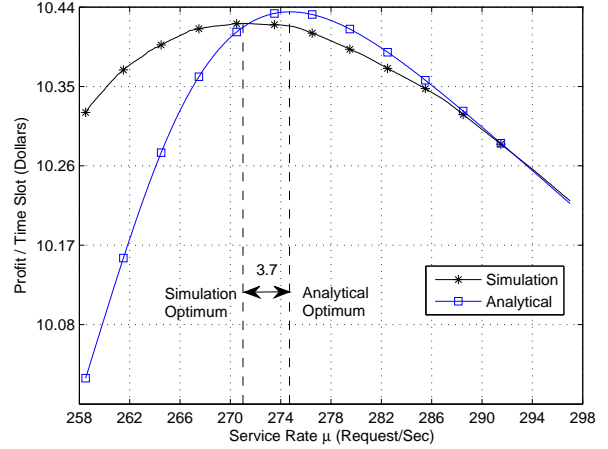


Fig. 3. Comparing analytical and simulation results in calculating data center profit as a function of μ over a sample $T = 15$ minutes time slot.

IV. PERFORMANCE EVALUATION

A. Simulation Setting

Consider a data center with a maximum of $M_{max} = 50,000$ servers. The exact number of switched on servers M is updated periodically at the beginning of each time slot of length $T = 15$ minutes by solving Problems (14) and (15) for the cases without and with behind-the-meter renewable power generation respectively. For each switched on server, we have $P_{peak} = 200$ watts and $P_{idle} = 100$ watts [10]. We assume that $E_{usage} = 1.2$ [12]. The electricity price information is based on the hourly real-time pricing tariffs currently practiced in Illinois Zone I, spanning from June 10, 2011 to July 9, 2011 [29]. We also assume that $\kappa = 0.1$ and consider the SLA to be according to the Gold SLA curve of Fig. 2. To simulate the total workload, we use the publicly available World Cup 98 web hits data, spanning from June 10, 1998 to July 9, 1998, as the trend for rate of incoming service requests [30].

B. Simulation Results for a Single Time Slot

To gain insights about the achievable profit, in this section, we focus on a single time slot of length $T = 15$ minutes at 1:00 PM on June 19 and investigate the solution of the profit maximization problem in (14). In Fig. 3, the analytical profit curve is compared with the profit curve that is calculated using an event-based simulator. We can see that the analytical curve is a close approximation of the simulation one. Based on the analytical curve, the optimum μ is 274.7 requests/second. This is only 3.7 requests/second greater than the true optimum service rate obtained from the simulation curve. That is, the optimality loss due to analytical modeling error is only 1.4%.

C. Simulation Results for an Entire Day

Fig. 4 shows a comparison between the proposed design with the one in [11] and [20] by simulation over 24 hours operation of the data center based on the June 19 data. The time-of-day price of electricity [29] as well as the workload [30] are shown in Fig. 4(e) and Fig. 4(d), respectively. We can

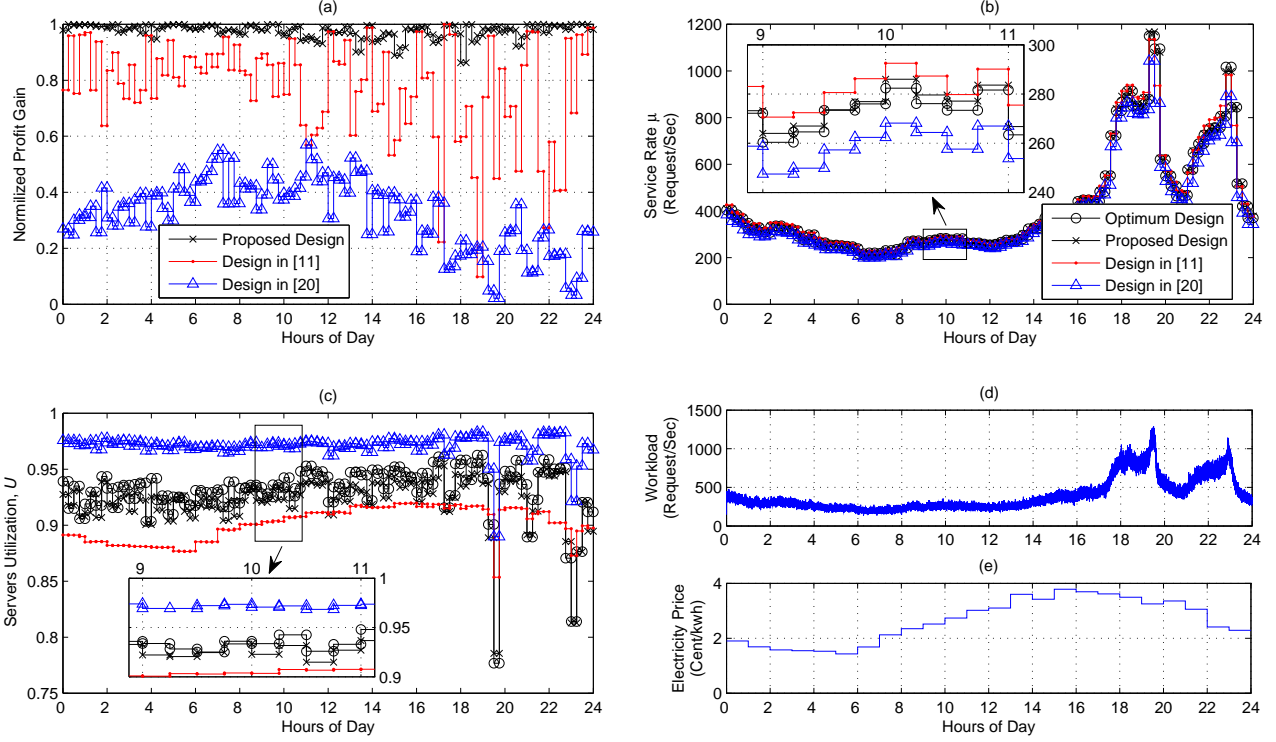


Fig. 4. Performance comparison between the proposed design and the designs in [11] and [20] over 24 hours running time of a data center: (a) Normalized profit gain. (b) The choice of service rate μ . (c) Servers' utilization. (d) Total workload [30]. (e) Hourly price of electricity [29].

see that the normalized profit gain of the proposed design in Fig. 4(a) is very close to optimal. The normalized profit gain is calculated as $(Profit - Profit_{Base})$ divided by $(Profit_{Max} - Profit_{Base})$. Here, $Profit_{Base}$ is the profit obtained when we simply set $\mu = \lambda$ [31] and $Profit_{Max}$ is the maximum of the profit curve obtained by simulation. Furthermore, we can see that our proposed design can outperform the two designs in [11] and [20]. The reason is two-fold. First, our design explicitly takes into account a mathematical model for data center's profit as a function of service rate. Second, we use an accurate G/D/1 queuing model while the designs in [11] and [20] are based on less accurate M/M/1 queuing models which cannot capture the workload distribution well. Finally, Fig. 4(b) shows the optimum choice of service rate, obtained from various designs, and Fig. 4(c) shows the server utilization U for the case of each design. We can see that the proposed design can achieve close to optimal service rates as well as close to optimal server utilization levels. The designs in [11] and [20] result in *under* and *over* utilized servers, respectively.

Next, we compare the normalized daily profit gain achieved by our design with those obtained by the designs in [11] and [20] over 30 days operation of the data center. The results are shown in Fig. 5. We can see that the proposed design works better in all days, with an average optimality of 96.3%. It is interesting to also point out that in 84% of the total $30 \times 24 \times 4 = 2880$ time slots being simulated, the optimum analytical service rate, i.e., the maximizer of the analytical profit curve, drops within interval (16), which is the range in which Problem (14) is a convex program. Finally, the normalized profit gain

of the three aforementioned designs, when $S = 1/\kappa$ changes, are compared in Fig. 6 for June 19 data. We can see that the proposed design is not sensitive to the service time parameter and can outperform [11] and [20] in all cases.

The loss probability of dropping a service request for different design approaches over 24 hours running time of the data center are shown in Fig. 7. Here, the choice of simulation parameters are the same as those in Fig. 4. The results in Fig. 7 show that the loss probability for our proposed design is closer to the optimum loss probability, compared to the cases of the designs in [11] and [20]. In order to understand the cause for this observation, we note that based on the results in Fig. 4-(c) the designs in [11] and [20] under-utilize and over-utilize the servers, respectively. For an under-utilized design, the loss probability is lower than the optimum value. For an over utilized design, the loss probability is higher than the optimum value. Clearly, we can draw a similar conclusion based on the loss probability results shown in Fig. 7.

D. Impact of Behind-the-Meter Renewable Generation

Next, assume that the data center is equipped with one local 1.5 Megawatts wind turbine. The power output for the turbine is assumed to be as in Fig. 8(a), based on the June 10 wind speed data available in [32]. In this case, the optimal service rate is obtained by solving Problem (15). The corresponding additional profit (in percentage) due to local renewable generation is shown in Fig. 8(b). We can see that local renewable generation can significantly increase the data center's daily profit by offsetting part of its energy cost.

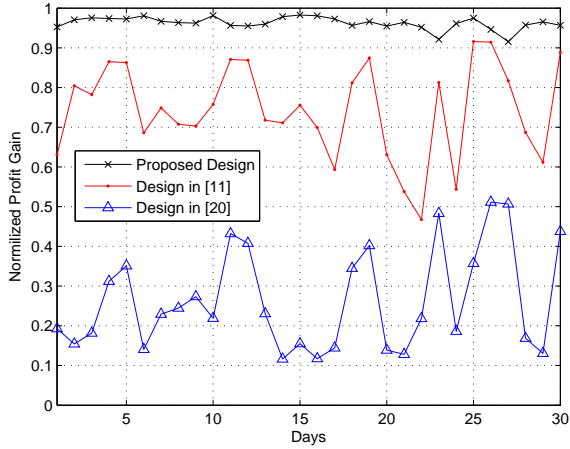


Fig. 5. Daily normalized profit gain across 30 days.

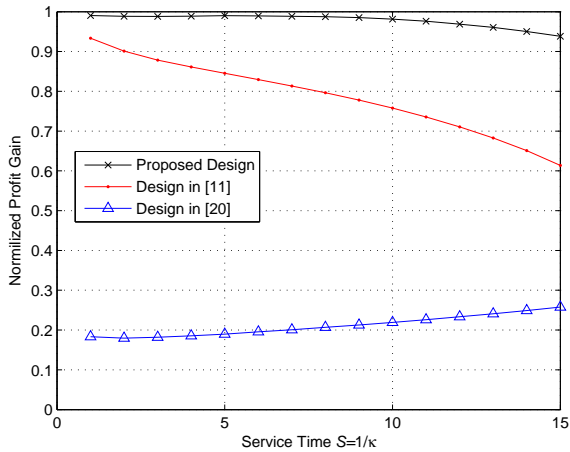


Fig. 6. Normalized profit gain as a function of $S = 1/\kappa$.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a novel analytical model to calculate profit in large data centers without and with behind-the-meter renewable power generation. Our model takes into account several factors including the practical service-level agreements that currently exist between data centers and their customers, price of electricity, and the amount of renewable power available. We then used the derived profit model to develop an optimization-based profit maximization strategy for data centers. We showed, under certain practical assumptions, the formulated optimization problems are convex. Finally, using various experimental data and via computer simulations, we assess the accuracy of the proposed mathematical model for profit and also the performance of the proposed optimization-based profit maximization strategy.

The results in this paper can be extended in several directions. First, the considered cost model can be generalized to include cost elements other than energy cost. Second, given renewable energy forecasting models and day-ahead pricing tariffs, the proposed short-term energy and performance management can be further extended to daily or monthly planning. Finally, the obtained mathematical model in this paper can

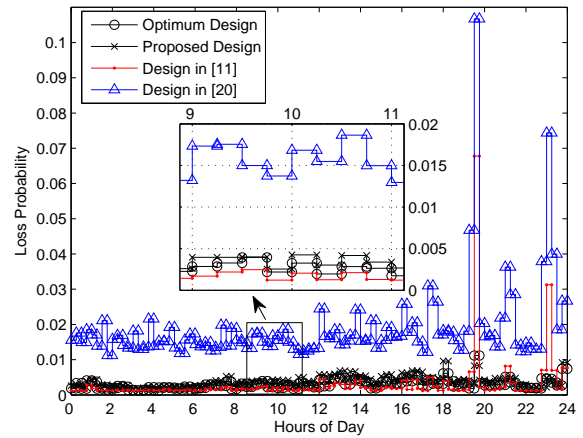


Fig. 7. Probability of dropping a service request due to full queues for various design approaches over 24 hours operation of data center.

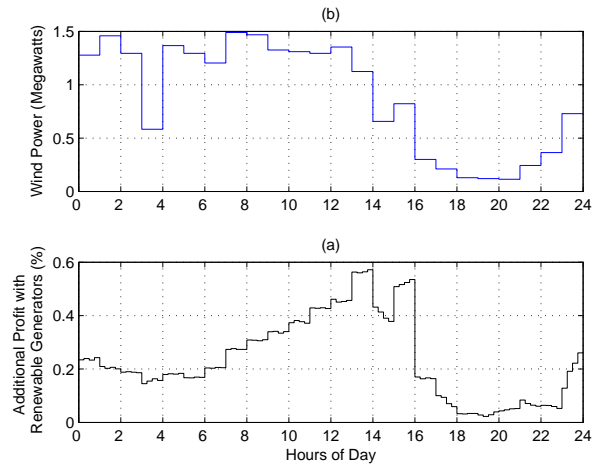


Fig. 8. Additional profit gained in a sample 24 hours operation of a data center with local renewable generation. Wind data is obtained from [32].

be adjusted to also include potential profit if a data center participates in ancillary services market in the smart grid.

REFERENCES

- [1] R. Katz, "Tech titans building boom," *IEEE Spectrum*, vol. 46, no. 2, pp. 40–54, Feb. 2009.
- [2] H. Mohsenian-Rad and A. Leon-Garcia, "Coordination of cloud computing and smart power grids," in *Proc. of the IEEE Smart Grid Communications Conference*, Gaithersburg, MD, oct 2010.
- [3] M. K. S. Gurumurthi, A. Sivasubramaniam and H. Franke, "Reducing disk power consumption in servers with drpm," *Computer*, vol. 36, no. 12, pp. 59–66, Dec. 2003.
- [4] J. Heo, D. Henriksson, X. Liu, and T. Abdelzaher, "Integrating adaptive components: An emerging challenge in performance-adaptive systems and a server farm case-study," in *Proc. of the IEEE International Real-Time Systems Symposium*, Tucson, AZ, Dec. 2007.
- [5] E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath, "Dynamic cluster reconfiguration for power and performance," in *Compilers and Operating Systems for Low Power*, M. K. L. Benini and J. Ramanujam, Eds. Kluwer Academic Publishers, 2003.
- [6] X. Fan, W. D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *Proc. of the ACM International Symposium on Computer Architecture*, San Diego, CA, June 2007.
- [7] J. S. Chase, D. Anderson, P. Thakar, A. Vahdat, and R. Doyle, "Managing energy and server resources in hosting centers," in *Proc. of ACM Symp. on Operating systems principles*, Banff, Canada, Oct. 2001.

- [8] L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing electricity cost: Optimization of distributed internet data centers in a multi-electricity-market environment," in *Proc. of IEEE INFOCOM*, Orlando, FL, 2010.
- [9] J. Li, Z. Li, K. Ren, and X. Liu, "Towards optimal electric demand management for internet data centers," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 183–192, Mar. 2012.
- [10] M. Ghamkhari and H. Mohsenian-Rad, "Optimal integration of renewable energy resources in data centers with behind-the-meter renewable generators," in *Proc. of the IEEE International Conference on Communications*, Ottawa, Canada, June 2012.
- [11] Z. Liu, M. Lin, A. Wierman, S. Low, and L. L. H. Andrew, "Geographical load balancing with renewables," in *Proc. of the ACM GreenMetrics Workshop*, San Jose, CA, Apr. 2011.
- [12] United States Environmental Protection Agency, "EPA report on server and data center energy efficiency," Final Report to Congress, Aug. 2007.
- [13] A. Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, "Cutting the electric bill for internet-scale systems," in *Proc. of the ACM SIGCOMM*, Barcelona, Spain, 2009.
- [14] P. R. S. Rivoire and C. Kozyrakis, "A comparison of high-level full-system power models," in *Proceedings of the 2008 conference on Power aware computing and systems*, ser. HotPower'08, 2008.
- [15] T. W. S. Pelley, D. Meisner and J. VanGilder, "Understanding and abstracting total data center power," in *Proc. of the ACM Workshop on Energy Efficient Design*, ser. HotPower'08, June 2009.
- [16] M. Ghamkhari and H. Mohsenian-Rad, "Data centers to offer ancillary services," in *Proc. of the IEEE Conference on Smart Grid Communications (SmartGridComm12)*, October 2012.
- [17] P. Torcellini, S. Pless, and M. Deru, "Zero energy buildings : A critical look at the definition preprint," *Energy*, p. 15, 2006.
- [18] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang, "Power and performance management of virtualized computing environments via lookahead control," in *International Conference on Autonomic Computing*, June 2008, pp. 3–12.
- [19] H. S. Kim and N. B. Shroff, "Loss probability calculations and asymptotic analysis for finite buffer multiplexers," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 755–768, Dec 2001.
- [20] L. Rao, X. Liu, L. Xie, and W. Liu, "Coordinated energy cost management of distributed internet data centers in smart grid," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 50–58, March 2012.
- [21] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowstron, "Better never than late: Meeting deadlines in datacenter networks," in *Proc. of the ACM SIGCOMM*, Aug. 2012.
- [22] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "Controlled experiments on the web: survey and practical guide," *Data Mining and Knowledge Discovery*, vol. 18, no. 1, pp. 140–181, Feb. 2009.
- [23] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: Amazons highly available key-value store," in *Proc. of the ACM SIGOPS*, Oct. 2007.
- [24] C. Y. Hong, M. Caesar, and P. B. Godfrey, "Finishing flows quickly with preemptive scheduling," in *Proc. of the ACM SIGCOMM*, Aug. 2012.
- [25] J. Rubio-Loyola, A. Galis, A. Astorga, J. Serrat, L. Lefevre, A. Fischer, A. Paler, and H. Meer, "Scalable service deployment on software-defined networks," *IEEE Communications Magazine*, vol. 12, no. 49, pp. 84–93, Dec. 2011.
- [26] D. Zats, T. Das, and R. H. Katz, "DeTail: Reducing the flow completion time tail in datacenter networks," Technical report, EECS Department, University of California, Berkeley, Mar. 2012.
- [27] T. Wood, P. Shenoy, A. Venkataramani, and M. Younis, "Sandpiper: Black-box and gray-box resource management for virtual machines," *Computer Networks*, vol. 53, no. 17, p. 29232938, 2009.
- [28] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.
- [29] <https://www2.amerem.com/RetailEnergy/realtimereprices.aspx>.
- [30] [Http://ita.ee.lbl.gov/html/contrib/WorldCup.html](http://ita.ee.lbl.gov/html/contrib/WorldCup.html).
- [31] X. Zheng and Y. Cai, "Reducing electricity and network cost for online service providers in geographically located internet data centers," in *IEEE/ACM International Conference on Green Computing and Communications*, Aug. 2011, pp. 166–169.
- [32] <http://www.windenergy.org/datasites/14-whitedeer/index.html>.
- [33] M. Abramowitz and I. A. Stegun, "Handbook of mathematical functions with formulas, graphs, and mathematical tables," Aug. 2007.

A. Proof of Theorem 1

First, we show that $q(\mu)$ in (10) is non-increasing. We define

$$t \triangleq (\mu - \lambda)/\sigma. \quad (21)$$

From (21), and after reordering the terms, we can show that

$$\alpha(\mu) = \frac{\sigma}{\lambda\sqrt{2\pi}} \left(1 - te^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du \right). \quad (22)$$

Once we take the derivative with respect to μ , we have

$$\alpha'(\mu) = \frac{1}{\lambda\sqrt{2\pi}} \left(t - (t^2 + 1)e^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du \right). \quad (23)$$

From [33, Formula 7.1.13], the following bounds always hold:

$$\frac{2}{t + \sqrt{(t^2 + 4)}} < e^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du \leq \frac{2}{t + \sqrt{(t^2 + 8/\pi)}}. \quad (24)$$

From the lower bound in (24) and the trivial inequality

$$\frac{t}{t^2 + 1} \leq \frac{2}{t + \sqrt{t^2 + 4}}, \quad (25)$$

we have $\alpha'(\mu) \leq 0$. That is, $\alpha(\mu)$ is non-increasing. On the other hand, for each $n \geq 1$, we have $m'_n(\mu) \geq 0$ [19]. As a result, $\exp(-\frac{1}{2} \min_{n \geq 1} m_n(\mu))$ is non-increasing. Therefore, $q(\mu)$ in (10) is non-increasing, i.e., $q'(\mu) \leq 0$.

Next, we prove that $q(\mu)$ is convex over interval (16). From (10), and since e^{-x} is non-increasing, we have

$$q(\mu) = \max_{n \geq 1} \alpha(\mu) e^{-\frac{1}{2} m_n(\mu)}. \quad (26)$$

Therefore, from [28, Section 3.2.3], $q(\mu)$ is proven to be a convex function if we can show that for each $n \geq 1$, function

$$q_n(\mu) \triangleq \alpha(\mu) e^{-\frac{1}{2} m_n(\mu)} \quad (27)$$

is convex. That is, for each $n \geq 1$, the second derivative

$$q_n''(\mu) = e^{-\frac{1}{2} m_n(\mu)} \left(\alpha''(\mu) + \alpha(\mu) \frac{m_n''(\mu)}{4} - \alpha'(\mu) m_n'(\mu) - \alpha(\mu) \frac{m_n''(\mu)}{2} \right) \geq 0. \quad (28)$$

Next, we show (28) through the following five steps:

Step 1: We show that $\alpha(\mu) \geq 0$. First, we note that

$$(r - \mu) e^{-\frac{(r-\lambda)^2}{2\sigma^2}} \geq 0, \quad \forall r \in [\mu, \infty]. \quad (29)$$

Therefore, the integral in (11) is non-negative. From this, together with the fact that $1/\lambda\sqrt{2\pi}\sigma$ and $\exp(\frac{(\mu-\lambda)^2}{2\sigma^2})$ are both non-negative terms, we can readily conclude that $\alpha(\mu) \geq 0$.

Step 2: We show that $\alpha''(\mu) \geq 0$ over interval (16). After taking the second derivative of α respect to μ , we have

$$\alpha''(\mu) = \frac{1}{\lambda\sqrt{2\pi}\sigma} \left((t^2 + 2) - (t^3 + 3t)e^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du \right). \quad (30)$$

Using simple algebra, we can show that

$$\frac{2}{t + \sqrt{(t^2 + 8/\pi)}} \leq \frac{t^2 + 2}{t^3 + 3t} \quad (31)$$

for all

$$0 \leq t \leq \sqrt{\frac{16/\pi}{3 - \frac{8}{\pi} + \sqrt{9 - \frac{16}{\pi}}}}. \quad (32)$$

Note that, from (21), condition (32) holds since (16) holds. Together, from (30), (31), and the upper bound in (24), we can directly conclude that $\alpha''(\mu) \geq 0$ over interval (16).

Step 3: We show that

$$-\frac{\alpha'(\mu)}{\alpha(\mu)} \geq \frac{1}{2\sigma} \left(\sqrt{t^2 + 4} - t \right). \quad (33)$$

From (22) and (23), we have

$$-\frac{\alpha'(\mu)}{\alpha(\mu)} = \frac{1}{\sigma} \left(\frac{e^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du}{1 - te^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du} - t \right). \quad (34)$$

From the lower bound in (24) and after reordering, we have

$$1 - te^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du \leq 1 - \frac{2t}{\sqrt{t^2 + 4} + t}. \quad (35)$$

Together, from (35) and the lower bound in (24), we have

$$\frac{e^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du}{1 - te^{\frac{t^2}{2}} \int_t^\infty e^{-\frac{u^2}{2}} du} \geq \frac{\frac{2}{t + \sqrt{t^2 + 4}}}{1 - \frac{2t}{t + \sqrt{t^2 + 4}}} = \frac{2}{\sqrt{t^2 + 4} - t}. \quad (36)$$

By replacing (36) in (34) and after reordering we obtain (33).

Step 4: We show that, over interval (16), we have

$$\frac{m_n''(\mu)}{m_n'(\mu)} \leq \frac{1}{\sigma} \left(\sqrt{t^2 + 4} - t \right). \quad (37)$$

From (12) and after taking the derivatives over μ , we have:

$$\frac{m_n''(\mu)}{m_n'(\mu)} = \frac{D + n}{D\mu + n(\mu - \lambda)} \leq \frac{1}{(\mu - \lambda)} = \frac{1}{\sigma t}, \quad (38)$$

where the last equality is due to (21). Note that, for each $n \geq 1$, we always have $m_n(\mu) \geq 0$, $m_n'(\mu) \geq 0$, and $m_n''(\mu) \geq 0$ [19]. Next, from the lower bound in (16), we have $t \geq \frac{\sqrt{2}}{2}\sigma$. From this, and by applying simple algebra, we can show that

$$\frac{1}{\sigma t} \leq \frac{1}{\sigma} \left(\sqrt{t^2 + 4} - t \right). \quad (39)$$

Replacing (39) in (38), we can directly conclude (37).

Step 5: From Steps 3 and 4, over interval (16), we have

$$\begin{aligned} \frac{m_n''(\mu)}{m_n'(\mu)} &\leq -2 \frac{\alpha'(\mu)}{\alpha(\mu)} \Rightarrow \\ -\alpha'(\mu)m_n'(\mu) - \alpha(\mu) \frac{m_n''(\mu)}{2} &\geq 0. \end{aligned} \quad (40)$$

Furthermore, from Steps 1 and 2, over interval (16), we have

$$\alpha''(\mu) + \alpha(\mu) \frac{m_n'^2(\mu)}{4} \geq 0. \quad (41)$$

From (40) and (41) and since the exponential function is non-negative, we can conclude (28) and the proof is complete. ■

B. Proof of Theorem 2

Part (a): We can rewrite the objective function in (14) as

$$T\lambda(\delta + \omega b/\kappa) - (T\omega a/\kappa)\mu - T\lambda q(\mu)(\delta - \omega b/\kappa + \gamma). \quad (42)$$

From Theorem 1, $q(\mu)$ is convex. Therefore, we need to show that the coefficient of $q(\mu)$ in (42) is non-negative. We have

$$\delta - \omega b/\kappa + \gamma \geq \delta - \omega b/\kappa > 0, \quad (43)$$

where the strict inequality is due to (18).

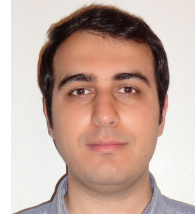
Part (b): From Theorem 1, $T\lambda [(1 - q(\mu))\delta - q(\mu)\gamma]$ and $(a\mu + b\lambda(1 - q(\mu)))/\kappa - G$ are both non-decreasing. Clearly, they are also continuous functions. Hence, for each μ , where

$$a\mu + b\lambda(1 - q(\mu)) < G, \quad (44)$$

there exists an $\epsilon > 0$, such that we have

$$a(\mu + \epsilon) + b\lambda(1 - q(\mu + \epsilon)) = G, \quad (45)$$

and the objective value of Problem (15) at service rate $\mu + \epsilon$ becomes no less than the objective value at service rate μ . Therefore, at optimality of Problem (15), we have $a\mu + b\lambda(1 - q(\mu)) \geq G$. This leads to the formulation of equivalent Problem (19). Next, we show that Problem (19) is convex. From Part (a), the objective function of Problem (19) is concave as it is the same as the objective function in Problem (14). Moreover, since $q(\mu)$ is convex, the expression $a\mu + b\lambda(1 - q(\mu))$ is a concave function in μ and consequently, the constraint in Problem (19) forms a convex set [28]. ■



Mahdi Ghamkhari (S'13) received his B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2011, and his M.Sc. degree from Texas Tech University, Lubbock, TX, USA, in 2012, both in Electrical Engineering. Currently, he is a PhD student in the Department of Electrical Engineering at the University of California, Riverside, CA, USA. His research interests include, Data Center and Cloud Computing, Convex and Stochastic Optimization, Queuing Theory and application of Game theory in deregulated Power Market. Mr. Mahdi Ghamkhari

has been a recipient of several graduate student awards both at Texas Tech University and at the University of California at Riverside.



Hamed Mohsenian-Rad (S'04-M'09) is currently an Assistant Professor of Electrical Engineering at the University of California at Riverside. He received his Ph.D. degree in Electrical and Computer Engineering from The University of British Columbia (UBC) - Vancouver in 2008. Dr. Mohsenian-Rad is the recipient of the NSF CAREER Award and the Best Paper Award from the IEEE International Conference on Smart Grid Communications 2012. He is an Associate Editor of the IEEE Communications Surveys and Tutorials, an Associate Editor of the

IEEE Communication Letters, a Guest Editor of the ACM Transactions on Embedded Computing Systems - Special Issue on Smart Grid, and a Guest Editor of the KICS/IEEE Journal of Communications and Networks - Special Issue of Smart Grid. His research interests include the design, optimization, and game-theoretic analysis of power systems and smart power grids.